

**METHODS AND APPARATUS FOR GENERATING
GRAPHICAL AND MEDIA DISPLAYS AT A CLIENT**

BACKGROUND

Cross-Reference to Related Application

[0001] This application claims priority to co-pending U.S. patent application number 10/098,157, filed March 14, 2002, the entirety of which is incorporated herein by reference.

Field Of Invention

[0002] The invention generally relates to distributed processing, and, more particularly, to generating a display having graphical and/or media components at a client.

Description of Prior Art

[0003] A thin-client protocol is used for displaying output, produced by an application running on a server, on a client with limited processing capabilities. Two exemplary thin client protocols are ICA, Independent Computing Architecture from Citrix Systems, Inc., Ft. Lauderdale, FL and RDP, Remote Desktop Protocol from Microsoft, Inc., Redmond, WA. The client is also sometimes referred to as a remote terminal session. One thin-client protocol intercepts commands by the application program to the server operating system ("OS") to draw to a display screen. The intercepted commands are transmitted to the remote session using, for example, one or more presentation layer packets. When the remote session (e.g., thin-client) receives the command, the remote session passes the received commands to the remote session OS. The thin-client draws the application program output on its display using the received commands. In this manner, the application program appears to be executing on the thin-client.

[0004] Typically, when the application program draws images to the display screen, the image is represented as a bitmap. A bitmap format of an image is generally a very large data set. Thus, the thin-client protocol must transmit over the network the bitmap representation of an image, which is a large amount of data, along with the applicable commands on how to display the bitmap representation. For networks of low bandwidth, this results in a large time delay before the complete image is received and displayed on the client. This can result in inconvenience and unhappiness for the user of the client. Also, if the user is paying directly for bandwidth used, for example in a wireless network, transmission of these large bitmap formats results in large costs associated with each transmission.

- 2 -

[0005] A similar problem exists when the application program renders a media presentation. Typically, a video file is rendered as a series of bitmaps and audio information is rendered using pulse code modulation. Accordingly, the thin-client protocol transmits the series of bitmaps representing the video file and/or the pulse code modulated signal representing the audio information over the network. This transmission is inefficient, requiring excessive bandwidth and significant CPU usage. Moreover, even where sufficient bandwidth is available, an unresponsive graphical user interface may result at the client. Video playback, for example, is often of low quality, may appear “jerky,” and may synchronize poorly with the audio presentation.

[0006] There is, therefore, a need for an improved approach to rendering images and media presentations in a remote terminal session.

SUMMARY OF THE INVENTION

[0007] The invention, according to one advantage, lowers the time and cost of transmitting images and other non-textual elements, originally represented in large bitmap formats, by substituting, prior to transmission, available compressed formats for the bitmap formats. Typically, transmitting the compressed formats instead of the bitmap formats can significantly reduce the necessary bandwidth. The client decompresses the received data using available libraries. The client then substitutes the decompressed image for the original bitmap representations using, for example, modified thin-client protocol commands with other identifying data.

[0008] According to another advantage of the invention, a compressed data set, representing at least a portion of a media stream, is intercepted on a first computing device before it is decompressed. Alternatively, where the compressed data set is decompressed on the first computing device, the resulting decompressed data set is re-compressed on the first computing device. By transmitting the compressed data set, rather than the decompressed data set, over a network, which may have limited bandwidth, the time and cost of transmitting the data set is consequently reduced.

[0009] In one aspect, the invention relates to a method for generating a graphical display at a client. The method includes transmitting output from an application program executing on a server to the client, identifying a bitmap representation within the application output, and determining a check value for the bitmap representation. The method also includes retrieving a compressed data format of the bitmap representation using at least in part the check value and transmitting to the client the compressed data format in place of the bitmap representation.

- 3 -

[0010] In another aspect, the invention relates to a method for generating a graphical display at a client. The method includes transmitting output from an application program executing on a server to the client and identifying a non-textual element within the application output. The method also includes retrieving a compressed data format associated with the non-textual element and transmitting to the client the compressed data format in place of the non-textual element.

[0011] In one embodiment of this aspect of the invention, the method includes identifying a textual element within the application output and transmitting to the client the textual element. In another embodiment, the method includes receiving the compressed data format, and optionally the textual element, at the client and generating a display at the client using the compressed data format, and optionally the textual element. In another embodiment, the method includes transmitting the compressed data format using at least one presentation layer protocol packet. In yet another embodiment, the method includes transmitting the at least one presentation layer protocol packet using a command for transmitting a file in its native format.

[0012] In another embodiment, the method includes conforming the at least one presentation layer protocol packet to a remote access protocol, a thin-client protocol, and/or a presentation protocol. In still another embodiment, the non-textual element is a bitmap representation and the method includes replacing the bitmap representation with the compressed data format. In another embodiment, the method includes determining the capability of the client to render the non-textual element using the compressed data format. The method further includes, upon determination that the client cannot render the non-textual element using the compressed data format, transmitting an image-rendering library capable of rendering the non-textual element using the compressed data format.

[0013] In another embodiment, the method includes intercepting the application output and inspecting the intercepted output for a bitmap representation of the non-textual element. In yet another embodiment, the method includes calculating a first check value for a bitmap representation of the non-textual element and searching an image store for the compressed data format having a check value identical to the first check value.

[0014] In another aspect, the invention relates to a system for generating a graphical display at a client. The system includes an output filter module and a server agent. The output filter module is configured to intercept output produced by an application program, identify a non-textual element of the output, and retrieve a compressed data format associated with the non-

- 4 -

textual element. The server agent is configured to transmit to the client the compressed data format in place of the non-textual element.

[0015] In one embodiment of this aspect of the invention, the system includes a server node, which includes the server agent and the output filter module. In another embodiment, the system includes a client node. The client node includes a client agent and a display. The client agent is configured to receive the compressed data format and to generate a display of the non-textual element using the received compressed data format. In another aspect, the system further includes a network.

[0016] In another aspect the invention relates to an article of manufacture having computer-readable program means embodied therein for generating a graphical display at a client. The article includes computer-readable program means for performing any of the aforementioned methods.

[0017] In an additional aspect, the invention relates to a method for generating a media presentation at a client. The method includes transmitting output from an application program executing on a server to the client, identifying a media stream within the application output, intercepting an original compressed data set representing at least a portion of the media stream before processing by the application program, and transmitting the original compressed data set to the client.

[0018] In another aspect, the invention relates to another method for generating a media presentation at a client. This method includes transmitting output from an application program executing on a server to the client, identifying a media stream within the application output, intercepting a first decompressed data set representing at least a portion of the media stream, compressing the intercepted first decompressed data set, and transmitting the compressed data set to the client in place of the first decompressed data set.

[0019] In yet another aspect, the invention relates to still another method for generating a media presentation at a client. This method includes informing a server of at least one media format supported by a client agent installed on the client, receiving a compressed data set of a media stream at the client, decompressing the compressed data set at the client to generate a decompressed data set, and generating the media presentation at the client using the decompressed data set.

[0020] In a further aspect, the invention relates to an article of manufacture that embodies computer-readable program means for generating a media presentation at a client. The article includes computer-readable program means for transmitting output from an application program

- 5 -

executing on a server to the client, computer-readable program means for identifying a media stream within the application output, computer-readable program means for intercepting an original compressed data set representing at least a portion of the media stream before processing by the application program, and computer-readable program means for transmitting the original compressed data set to the client.

[0021] In still another aspect, the invention relates to another article of manufacture that embodies computer-readable means for generating a media presentation at a client. This article includes computer-readable program means for transmitting output from an application program executing on a server to the client, computer-readable program means for identifying a media stream within the application output, computer-readable program means for intercepting a first decompressed data set representing at least a portion of the media stream, computer-readable program means for compressing the intercepted first decompressed data set, and computer-readable program means for transmitting the compressed data set to the client in place of the first decompressed data set.

[0022] In yet another aspect, the invention relates to yet another article of manufacture that embodies computer-readable means for generating a media presentation at a client. This article includes computer-readable program means for informing a server of at least one media format supported by a client agent installed on the client, computer-readable program means for receiving a compressed data set of a media stream at the client, computer-readable program means for decompressing the compressed data set at the client to generate a decompressed data set, and computer-readable program means for generating the media presentation at the client using the decompressed data set.

[0023] In various embodiments of these last six aspects of the invention, the methods further include, and the articles of manufacture further include computer-readable program means for, capturing timing information associated with the media stream, transmitting the timing information to the client, receiving the compressed data set and, optionally, the timing information at the client, decompressing the compressed data set at the client to generate a decompressed data set, and generating the media presentation at the client using the decompressed data set and, optionally, the timing information. In other embodiments of the last six aspects of the invention, the methods further include, and the articles of manufacture further include computer-readable program means for, transmitting non-media graphical information from the application output to the client, receiving the non-media graphical information at the

- 6 -

client, and generating the media presentation at the client using the decompressed data set and the non-media graphical information.

[0024] In an additional aspect, the invention relates to a system for generating a media presentation at a client. The system includes an application program and an output filter module. The application program is configured to identify a media stream within output produced by the application program. The output filter module is configured to intercept an original compressed data set representing at least a portion of the media stream before processing by the application program and transmit the original compressed data set to the client.

[0025] In another aspect, the invention relates to another system for generating a media presentation at a client. This system includes an application program and an output filter module. The application program is configured to identify a media stream within output produced by the application program. The output filter module is configured to intercept a first decompressed data set representing at least a portion of the media stream, compress the intercepted first decompressed data set of the media stream, and transmit the compressed data set in place of the first decompressed data set to the client.

[0026] In yet another aspect, the invention relates to another system for generating a media presentation at a client. This system includes a server and the client in communication with the server. The client includes a client agent configured to inform the server of at least one media format supported by the client agent, receive a compressed data set of a media stream, decompress the compressed data set at the client to generate a decompressed data set, and generate the media presentation using the decompressed data set.

[0027] In various embodiments of these last three aspects of the invention, the output filter module of the systems is further configured to capture timing information associated with the media stream and to transmit the timing information to the client. In various other embodiments of the last three aspects of the invention, the system further includes a client agent configured to receive the compressed data set and the optional timing information, decompress the compressed data set to generate a decompressed data set, and generate the media presentation using the decompressed data set and the optional timing information. In still other embodiments of the last three aspects of the invention, the client agent is further configured to receive non-media graphical information and to generate the media presentation at the client using the decompressed data set and the non-media graphical information.

[0028] In another aspect, the invention relates to another system for generating a media presentation at a client. This system includes a network, a server in communication with the

- 7 -

network, and the client in communication with the network. The server includes an application program and at least one output filter module. The application program is configured to identify a media stream within output produced by the application program. The output filter module is configured to intercept a compressed data set representing at least a portion of the media stream before processing by the application program, and transmit the compressed data set to the client. The client includes a client agent. The client agent is configured to inform the server of at least one media format supported by the client agent, receive the compressed data set, decompress the compressed data set at the client to generate a decompressed data set, and generate the media presentation at the client using the decompressed data set.

[0029] In a further aspect, the invention relates to an article of manufacture that embodies computer-readable program means for generating a media presentation at a client. The article includes computer-readable program means for intercepting an original compressed data set of a media stream, and computer-readable program means for transmitting the original compressed data set to the client using a thin client protocol such as ICA or RDP.

[0030] In another aspect, the invention relates to another article of manufacture that embodies computer-readable program means for generating a media presentation at a client. The article includes computer-readable program means for intercepting a decompressed data set of a media stream, computer-readable program means for compressing the intercepted decompressed data set, and computer-readable program means for transmitting the compressed data set to the client using a thin client protocol such as ICA or RDP.

BRIEF DESCRIPTION OF THE DRAWINGS

[0031] The above and further advantages of the invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

[0032] FIG. 1 is a block diagram of an illustrative embodiment of a system to generate a graphical display for a remote terminal session in accordance with the invention;

[0033] FIG. 2 is a flow diagram of an illustrative embodiment of a process to generate a graphical display for a remote terminal session in accordance with the invention;

[0034] FIG. 3 is a block diagram of an illustrative embodiment of a system for generating a media presentation at a client in accordance with the invention; and

[0035] FIGS. 4A, 4B, & 4C are a flow diagram of an illustrative embodiment of a method for generating a media presentation at a client in accordance with the invention.

- 8 -

DETAILED DESCRIPTION

[0036] In one aspect, the invention pertains to methods, systems, and articles of manufacture for generating a graphical display. A compressed data format, associated with a non-textual element or a bitmap representation of an image, is transmitted over a network from a server to a client, in place of the non-textual element or the bitmap representation, for subsequent display.

[0037] In broad overview, FIG. 1 illustrates a system 100 to generate a display for a remote terminal session that includes a first computing system ("client node") 105 in communication with a second computing system ("server node") 110 over a network 115. For example, the network 115 can be a local-area network (LAN), such as a company Intranet, or a wide area network (WAN), such as the Internet or the World Wide Web. A user of the client node 105 can be connected to the network 115 through a variety of connections including standard telephone lines, LAN or WAN links (e.g., T1, T3, 56kb, X.25), broadband connections (ISDN, Frame Relay, ATM), and wireless connections. The client node 105 includes a client transceiver 130 to establish communication with the network 115. The server node 110 includes a server transceiver 135 to establish communication with the network 115. The connections can be established using a variety of communication protocols (e.g., ICA, RDP, HTTP, TCP/IP, IPX, SPX, NetBIOS, Ethernet, RS232, and direct asynchronous connections).

[0038] The server node 110 can be any computing device capable of providing the requested services of the client node 105. Particularly, this includes generating and transmitting commands and data to the client node 105 that represent the output being produced by an application program 140 executing on the server 110. The server node 110 includes the server transceiver 135, the executing application program 140, a server agent 150, an output filter module 155 and an image store 160. The server agent 150 includes a module that interfaces with a client agent 175 and other components of the server node 110 to support the remote display and operability of the application program 140. The server agent module 150 and all modules mentioned throughout the specification are implemented as a software program and/or a hardware device (e.g., ASICs or FPGAs).

[0039] For clarity, all of these components are shown on server node 110. It is to be understood that the server node 110 can represent a single server or can represent several servers in communication with each other over the network 115 or another network (not shown). In multiple server embodiments, the functionality of the components can be distributed over the available servers. For example, in one embodiment with multiple servers, the transceiver 135, the application program 140, the server agent 150 and the output filter module 155 are on an

- 9 -

application server and the image store 160 is on a storage device, such as a disk in a RAID system.

[0040] The client node 105 can be any computing device (e.g., a personal computer, set top box, wireless mobile phone, handheld device, personal digital assistant, kiosk, etc.) used to provide a user interface to the application program 140 executing on the server node 110. The client node 105 includes the client transceiver 130, a display 145, a client agent 175 and a graphics library 180 (also referred to as an image-rendering library). The client agent 175 includes a module, implemented as a software program and/or a hardware device (e.g., an ASIC or an FPGA) that receives commands and data from the server node 110 and from a user (not shown) of the client node 105. The client agent 175 uses the received information when interacting with other components of the client node 105 (e.g., when directing the operating system to output data onto the display 145). The client agent 175 also transmits requests and data to the server node 110 in response to server-issued commands or user actions at the client node 105.

[0041] The server node 110 hosts one or more application programs 140 that can be accessed by the client nodes 105. Examples of such applications include word processing programs such as MICROSOFT Word and spreadsheet programs such as MICROSOFT Excel, both manufactured by Microsoft Corporation of Redmond, Washington. Other examples include financial reporting programs, customer registration programs, programs providing technical support information, customer database applications, and application set managers. Another example of an application program is Internet Explorer, manufactured by Microsoft Corporation of Redmond, Washington, and this program will be used as an exemplary application program 140 in the following discussion. It is understood that other application programs can be used.

[0042] During execution of the application program 140, for example Internet Explorer, the server node 110 communicates with the client node 105 over a transport mechanism. In one embodiment, the transport mechanism provides multiple virtual channels 185 through the network 115 so the server agent 150 can communicate with the client agent 175. One of the virtual channels 185 provides a protocol for transmitting graphical screen data from the server node 110 to the client node 105. The server 110 executes a protocol driver, in one embodiment as part of the server agent 150, that intercepts graphical display interface commands generated by the application program 140 and targeted at the server's operating system. The protocol driver translates the commands into a protocol packet suitable for transmission over the transport mechanism.

- 10 -

[0043] The application program 140, in this example Internet Explorer, executing on the server 110, retrieves a web page. As explained above, the application program 140 generates graphical display commands to the server operating system, as if it was going to display the output at the server node 110. The server agent 150 intercepts these commands and transmits them to the client agent 175. The client agent 175 issues the same or similar commands to the client operating system to generate output for the display 145 of the client node 105.

[0044] In one embodiment, a web page has both textual elements (e.g., titles, text, and ASCII characters) and non-textual elements (e.g., images, photos, icons, and splash screens) incorporated therein. The non-textual elements are sometimes transmitted to the Internet Explorer application program 140 from a web server (not shown) in a compressed data format (e.g., a file or a data stream), also referred to as the non-textual element's native format. Examples of compressed formats are JPEG, GIF, and PNG. The non-textual element represented in a compressed data format may be, for example, 20 kilobytes in size. That same non-textual element decompressed into its bitmap representation is, for example, 300 kilobytes in size.

[0045] The application program 140, when generating the display of the web page, retrieves, for example, a JPEG data format of a non-textual element and decompresses the JPEG data format into a bitmap for display. The output filter module 155 determines that the bitmap representation is from a compressed format and obtains the corresponding compressed format of the non-textual element from the image store 160, as explained in more detail below. In one embodiment, the image store 160 is persistent storage. In other embodiments, the image store 160 is temporary storage, cache, volatile memory and/or a combination of temporary and persistent storage.

[0046] The server agent 150 replaces the bitmap representation of the non-textual element with the compressed non-textual element that the output filter module 155 retrieved from the image store 160. The server agent 150 transmits the non-textual element in the compressed format, along with the graphical display interface commands associated with the bitmap representation, to the client node 105. In one embodiment the server agent 150 uses a unique protocol command that identifies a transmission of a non-textual element that is not in bitmap representation, even though the associated commands are applicable to a bitmap representation of a non-textual element. In other embodiments other identifying techniques can be used. For example, the protocol command can have a modifier comment, or a command switch. The command can also use a change of context or a combination of multiple commands.

- 11 -

[0047] The client agent 175 receives the transmission of the non-textual element file in the compressed data format, along with the graphical display interface commands associated with the bitmap representation of the non-textual element. The client agent 175 determines that the non-textual element is in the compressed data format and not the bitmap representation. In one embodiment, the client agent 175 makes this determination because the non-textual element in compressed format is transmitted using a unique protocol command. In another embodiment, the size of the non-textual element data and/or other characteristics about the non-textual element included with the associated graphical display interface commands are enough to enable the client agent 175 to make the determination.

[0048] The client agent 175 determines whether the client node 105 contains the necessary library 180 to decompress the compressed format of the non-textual element. If the client node 105 has the appropriate graphics library(ies) 180 installed to perform the decompression algorithms, the client agent 175 uses the library 180 to decompress the compressed format of the non-textual element into its bitmap representation. The client agent 175 performs the received associated graphical display interface commands on the bitmap representation to generate the non-textual element of the output of the application program 140 on the client display 145.

[0049] In one embodiment, the client agent 175 does not contain all the decompression algorithms to decompress the non-textual element from a compressed format into a bitmap representation. If the client node 105 does not have the appropriate graphics library(ies) 180 installed to perform the decompression algorithms, the client agent 175 requests the needed graphics library from the server node 110. In another embodiment, the client agent 175 determines if a predetermined set of the most widely used graphics libraries 180 are installed on the client node 105 prior to receiving any non-textual elements from the server node 110. If the most widely used graphics libraries 180 are not installed on the client node 105, the client agent 175 requests the missing libraries from the server node 110 prior to receiving any non-textual elements from the server node 110.

[0050] In yet another embodiment, the client agent 175 determines which graphics libraries 180 the client node 105 includes and transmits that library information to the server agent 150. In this embodiment, when the server agent 150 receives the compressed data format of a bitmap representation from the output filter module 155, the server agent 150 determines, using the transmitted library information, whether the client node 105 can render the compressed data format. If the server agent 150 determines that the client node 105 has the applicable library, the server agent 150 substitutes the compressed data format for the bitmap representation of the non-

- 12 -

textual element. If the server agent 150 determines that the client node 105 does not have the applicable library, the server agent 150 does not substitute the compressed data format for the bitmap representation of the non-textual element and instead transmits the bitmap representation to the client 105.

[0051] For the server agent 150 to replace the bitmap representation of the non-textual element with the non-textual element in the compressed format, the output filter module 155 determines that the bitmap representation is from a compressed format contained in the image store 160. To make this determination, the output filter module 155 calculates one or more check values for the bitmap representation. For example, the output filter module 155 can calculate a single check value for the entire bitmap representation and/or the output filter module 155 can calculate four check values, one for each quadrant for the entire bitmap representation. In another example, the output filter module 155 can calculate N check values, one for each of the N lines in the bitmap representation. A check value is the result of an algorithm that generates a substantially unique value for different arrangements of data. The check value is, for example, a checktag, a Cyclic Redundancy Code ("CRC"), a check sum, or a result of a hashing function. The check value is based on the bitmap representation and not the data as arranged in a compressed data format. However, when the compressed data format is stored in the image store 160, it is stored with a check value attribute that corresponds to the one or more check values of the bitmap representation of the compressed data when decompressed.

[0052] In one embodiment, the check value is a checktag that includes a fixed identifier and a unique identifier. The fixed identifier and the unique identifier are combined together and concealed within an image. The fixed identifier is used to identify the checktag as such; the unique identifier is used to identify a specific image. The fixed identifier is, for example, a globally unique identifier that is statistically unlikely to be found within a image. For example, the fixed identifier is a byte sequence that is easily recognizable during debugging and that has a balanced number of zero and one bits. The unique identifier is a sequential identifier uniquely allocated for each image in the cache. The sequential unique identifier is XOR masked with another value so that the image identifiers with a small value (the most likely value) will be more likely to have a balanced number of zero and one bits.

[0053] The checktag is encoded into RGB color components, independently of whether the RGB components are part of the image or part of the color palette. More specifically, the checktag is treated as a stream of 160 bits (*i.e.*, 20 separate bytes, each of which starts at bit 0, the least significant, and finishes at bit 7, the most significant bit). The least significant bit of

- 13 -

each byte is overwritten by the next bit of the checktag. The other 7 bits of each byte remain unaltered.

[0054] A checktag is decoded by simply reversing the encoding procedure. After the checktag is decoded, the fixed identifier and the unique identifier are retrieved from the checktag. The retrieved fixed identifier is validated against a previously stored fixed identifier to identify the checktag as such. Where a match is found, the unique identifier is then used to retrieve information that is relevant to the identified image, such as the bitmap data associated with the image.

[0055] The output filter module 155 searches the image store 160 for a non-textual element in compressed data format that has a check value attribute that is the same as one or more check values the output filter module 155 calculates for the bitmap representation. The output filter module 155 retrieves the compressed format of non-textual element with the same check value attribute as the one or more check values and sends the compressed format of the non-textual element to the server agent 150 for transmittal to the client agent 175 in place of the bitmap representation.

[0056] The server node 110 stores compressed formats of non-textual elements in the image store 160 the first time the application program 140 calls a graphics library (not shown) to create a bitmap representation from a compressed format file. The output filter module 155 calculates the associated check value of the bitmap representation as the application program 140 decompresses the compressed format and generates the bitmap representation. As described above, the output filter module 155 can calculate the check value when the bitmap representation is complete, when a quadrant of the bitmap representation is complete, or when a line of the bitmap representation is complete. The server 110 stores the compressed format file and the associated check value attribute in the image store 160 and retrieves the compressed format file the first and any subsequent times the application program 140 generates the associated non-textual element.

[0057] Whether the server 110 stores the compressed format file and its associated check value attribute(s) in the image store 160 in a temporary portion (e.g., RAM memory buffer or cache) or a persistent portion (e.g., disk or non-volatile memory buffer) is based at least in part on design and hardware limitations (e.g., the size of the persistent storage). One exemplary criterion used to make that determination is the number of times the application program 140 generates the non-textual element. For example, if the application program 140 generates a particular non-textual element more than a predetermined number of times, the server 110 stores

- 14 -

the compressed format file and its associated check value attribute(s) corresponding to that particular non-textual element persistently in the image store 160.

[0058] In other embodiments, the server 110 stores the non-textual element if it is static or complex. For example, if the application program 140 always generates a splash screen at initialization, the server 110 stores the compressed format file corresponding to that splash screen in the persistent portion of the image store 160. In another embodiment, if the non-textual element is complex, static and/or generated repeatedly but does not have a corresponding compressed format file, the output filter module 155 generates a compressed format file for that non-textual element, in a standards-based or proprietary-based format. In any subsequent transmissions, the server agent 150 transmits the generated compressed format file in place of the bitmap representation. If the compressed format is a proprietary-based format, the server agent 150 determines whether the client node 105 includes the applicable proprietary-based graphics library to decompress the compressed format file into a bitmap representation. If not included on the client node 105, the server agent 150 transmits the applicable library to the client node 105 for installation.

[0059] Although the illustrated embodiment depicts the image store 160 on the server node 110, in an alternate embodiment, at least a portion of the image store (not shown) is on the client node 105. In this alternate embodiment, the output filter module 155 calculates the one or more check values of the bitmap representation and transmits the one or more check values to the server agent 150. The server agent 150 transmits these one or more check values to the client agent 175. The client agent 175 searches the portion of the image store on the client node 105 for a compressed data format stored with an identical one or more check values attribute. The client agent 175 transmits the results of this search to the server agent 150.

[0060] If the compressed data format for the non-textual element exists on the client node 105, the server agent 150 does not have to send either the compressed data format or the bitmap representation over the network 115. The server agent 150 only transmits the graphical display interface commands associated with the bitmap representation of the non-textual element. If the compressed data format for the non-textual element does not exist on the client node 105, the output filter module 155 obtains the corresponding compressed format of the non-textual element from the image store 160. The server agent 150 replaces the bitmap representation of the non-textual element with the non-textual element in the compressed data format that the output filter module 155 retrieved from the image store 160. The server agent 150 transmits the

- 15 -

non-textual element in the compressed format, along with the graphical display interface commands associated with the bitmap representation, to the client node 105.

[0061] FIG. 2 illustrates an exemplary embodiment of a process 200 to generate a display for a remote terminal session, using the exemplary embodiment of FIG. 1. The output filter module 155 monitors the output of the application program 140 by monitoring calls made to the operating system of the server node 110. When the output filter module 155 detects (step 205) a display command from the application program 140, the output module 155 determines (step 210) whether the application program 140 is generating a bitmap representation of a non-textual element.

[0062] If the application program 140 is not generating a bitmap representation of a non-textual element, the output filter module 155 transmits (step 215) the display command to the server agent 150, which transmits that command, or a representative command defined in the protocol, to the client agent 175. If the application program 140 is generating a bitmap representation of a non-textual element, the output filter module 155 calculates (step 220) one or more check values corresponding to the bitmap representation of the non-textual image.

[0063] Using the one or more calculated check value(s), the output filter module 155 searches the image store 160 to determine (step 225) whether a compressed data format with identical check value attribute(s) exists. If there is a compressed data format in the image store 160 with check value attribute(s) identical to the one or more check values the output filter module 155 calculates, the output module 155 replaces (step 230) the bitmap representation of the non-textual element with the compressed data format. The output module 155 transmits (step 230) the compressed data format to the server agent 150 for eventual transmission to the client agent 175. The output module 155 also transmits all of the commands associated with the replaced bitmap representation along with the compressed data format.

[0064] If there is not a compressed data format with identical one or more check value attributes in the image store 160, the output module 155 determines (step 235) whether the bitmap representation of a non-textual element corresponding to the compressed data format meets a predetermined criterion for persistent storage (e.g., any of the criteria described above). If the output module 155 determines (step 235) that the predetermined criterion is met, the output module 155 stores (step 240) the compressed data format and the corresponding check value attribute, identical to the one or more calculated check values, in the persistent portion of the image store 160. If the output module 155 determines (step 235) that the predetermined criterion is not met, the output module 155 stores (step 245) the compressed data format and the

- 16 -

corresponding check value attribute, identical to the one or more calculated check values, in the temporary portion of the image store 160.

[0065] Once the output module 155 stores (step 240 or 245) the compressed data format and the corresponding check value attribute, identical to the one or more calculated check values, in the image store 160, the output module 155 replaces (step 230) the bitmap representation of the non-textual element with the compressed data format. The output module 155 transmits (step 230) the compressed data format to the server agent 150 for eventual transmission to the client agent 175. The output module 155 continues monitoring the output generated by the application program 140 until the output module 155 detects (step 205) another display command from the application program 140.

[0066] In another aspect, the invention pertains to methods, systems, and articles of manufacture for generating a media presentation. In one embodiment, a compressed data set, representing at least a portion of a media stream, is intercepted on a first computing device and transmitted, over a network, to a second computing device, where it is decompressed and presented to a user. In another embodiment, a decompressed data set, representing at least a portion of a media stream, is intercepted and compressed on the first computing device and then transmitted, as above, to the second computing device, where it is again decompressed and presented to the user.

[0067] FIG. 3 illustrates one embodiment of a system 300 that generates a media presentation according to this aspect of the invention. The system 300 includes a first computing device, *e.g.*, a server 310, in communication with a second computing device, *e.g.*, a client 305, over a network 315. Generally speaking, except as set forth below, the client 305, the server 310, and the network 315 have the same capabilities as the client 105, the server 110, and the network 115, respectively, described above.

[0068] As shown, the client 305 includes at least a client transceiver 330, a client agent 375, and a presentation interface 345. The client agent 375 may be implemented as a software program and/or as a hardware device, such as, for example, an ASIC or an FPGA. The client agent 375 uses the client transceiver 330 to communicate over the network 315 and generates a presentation having media and non-media components at the presentation interface 345.

[0069] In one embodiment, the server 310 is an application server. As illustrated, the server 310 includes at least a server transceiver 335, an application program 340, a server agent 350, a first output filter module 355A, and a second output filter module 355B. The server agent 350, the first output filter module 355A, and the second output filter module 355B may be

- 17 -

implemented as a software program and/or as a hardware device, such as, for example, an ASIC or an FPGA. The server agent 350, the first output filter module 355A, and the second output filter module 355B use the server transceiver 335 to communicate over the network 315.

[0070] In another embodiment, the aforementioned components 335, 340, 350, 355A, and 355B are distributed over several servers in communication with each other over the network 315, or over another network (not shown). Alternatively, in yet another embodiment, two or more of the aforementioned components 335, 350, 355A, and 355B may be combined into a single component, such that the functions, as described below, performed by two or more of the components 335, 350, 355A, and 355B are performed by the single component.

[0071] The application program 340 illustrated in FIG. 3 is any application program 340 that renders, as part of its output, a media stream. The media stream may be a video stream, an audio stream, or, alternatively, a combination thereof. In addition, the application program 340 may output non-media graphical information. In this context, non-media graphical information refers generally to all graphical information outputted by the application program 340 without the use of a codec or the equivalent, such as, for example, static graphical information, including, but not limited to, toolbars and drop-down menus. Non-media graphical information also includes, for example, information for locating the static graphical information on a display screen. The application program 340 may be, for example, the MICROSOFT ENCARTA application program manufactured by the Microsoft Corporation of Redmond, Washington.

[0072] In one embodiment, the application program 340 uses external codecs, such as, for example, codecs installed in the operating system of the server 310, to decompress a compressed data set representing at least a portion of a media stream. In another embodiment, the codecs used by the application program 340 are embedded in the application program 340 itself. In yet another embodiment, the server 310 may include any number of executing application programs 340, some of which use external codecs, others of which use embedded codecs.

[0073] Where the application program 340 uses external codecs and desires to output a media stream, it requests that the operating system of the server 310 use the external codecs to decompress the compressed data set representing at least a portion of the media stream for subsequent display. Where the codecs used by the application program 340 are embedded in the application program 340 itself, the application program 340, when desiring to output a media stream, uses the embedded codecs to decompress the compressed data set itself for subsequent display. Additionally, the application program 340 may generate and transmit graphical display

- 18 -

commands, associated with the non-media graphical information, to the operating system of the server 310.

[0074] In accordance with the present invention, the application program 340 performs these tasks as if the application program 340 was going to generate a presentation having media and non-media components at the server 310. As explained below, the first output filter module 355A, the second output filter module 355B, and the server agent 350 intercept the compressed data set being passed to the external codecs, the decompressed data set generated by the embedded codecs, and the graphical display commands associated with the non-media graphical information, respectively, and (after first compressing the decompressed data set generated by the embedded codecs) transmit them, over the network 315, to the client agent 375. The client agent 375, as explained below, then decompresses the received compressed data sets and issues the same or similar graphical display commands, associated with the non-media graphical information, to the operating system of the client 305 to generate a presentation having media and non-media components at the presentation interface 345 of the client 305.

[0075] The first output filter module 355A and the second output filter module 355B are invoked as an application program 340 using external codecs attempts to invoke an external codec to output a media stream. The first output filter module 355A intercepts an original compressed data set representing at least a portion of the media stream. Instead of decompressing the data set, as an external codec would, the first output filter module 355A transmits the original compressed data set over the network 315 to the client agent 375. Where the media stream includes a video stream, the second output filter module 355B acting as an OS-level renderer captures information for locating images of the video stream on a display screen and transmits the information over the network 315 to the client agent 375.

[0076] In another embodiment, when an application program 340 that uses embedded codecs attempts to invoke, for example, an OS-level renderer to output a media stream, the second output filter module 355B is invoked. The second output filter module 355B intercepts a first decompressed data set representing at least a portion of the media stream from the output of the application program 340. The second output filter module 355B then compresses, as explained below, the intercepted first decompressed data set and transmits the resulting compressed data set, over the network 315, to the client agent 375. The second output filter module 355B, as above, also captures, where the media stream includes a video stream, information for locating images of the video stream on a display screen and transmits the information over the network 315 to the client agent 375.

- 19 -

[0077] In another embodiment, the server agent 350 intercepts and transmits to the client agent 375, over the network 315, the graphical display commands associated with the non-media graphical information outputted by the application program 340.

[0078] Additionally, in one embodiment, the first output filter module 355A, the second output filter module 355B or both (where the application program 340 uses external codecs), or the second output filter module 355B (where the application program 340 uses embedded codecs), captures timing information associated with the media stream and transmits the timing information, over the network 315, to the client agent 375. More specifically, the output filter module 355A, 355B captures, and transmits to the client agent 375, presentation times for each frame of the media stream, thereby enabling the client agent 375 to synchronize video and audio streams and to maintain the correct frame rate.

[0079] As shown, the server agent 350 interfaces with the server transceiver 335 and the application program 340. In one embodiment, as explained below, the server agent 350 receives from the client agent 375, over the network 315, a list of media formats supported by the client agent 375. Upon receiving the list of supported media formats, the server agent 350 registers the output filter modules 355A, 355B by manipulating the configuration of the server 310. In one embodiment, for example, the server agent 350 registers the output filter modules 355A, 355B by editing the registry of the server 310. The server agent 350 then informs the client agent 375 that the server 310 can handle all such media formats.

[0080] At the client 305, the client agent 375 interfaces with the client transceiver 330 and the presentation interface 345. The client agent 375, as described below, initially informs the server agent 350 of the media formats supported by the client agent 375. The client agent 375 also receives from the output filter modules 355A, 355B, over the network 315, the compressed data set and any associated timing information. Moreover, the client agent 375 receives over the network 315, from the second output filter module 355B, any information for locating images of a video stream on a display screen and, from the server agent 350, the graphical display commands associated with the non-media graphical information.

[0081] The client agent 375, using either external or embedded codecs, decompresses the compressed data set and, together with the graphical display commands associated with the non-media graphical information, any information for locating images of a video stream on a display screen, and any timing information, generates a media presentation at the presentation interface 345. The presentation interface 345 has, in one embodiment, a display screen that renders a graphical display, such as, for example, a video presentation. In another embodiment, the

- 20 -

presentation interface 345 includes a speaker that renders an audio presentation. The client 305 may include any number of presentation interfaces 345.

[0082] The information provided in specifying the media formats supported by the client agent 375 may determine the mode of operation at the server 310. If the compressed data set is in a format that is not supported by the client agent 375, the second output filter module 355B may recompress the decompressed data set into a supported format.

[0083] Referring now to FIGS. 4A, 4B, and 4C, one embodiment of a method 400 that generates a media presentation at the client 305, using the exemplary embodiment of FIG. 3, is illustrated. When the client 305 and the server 310 are both connected to the network 315, the client agent 375, at step 404, informs the server agent 350 of all the media formats supported by the client agent 375. In one embodiment, the list of supported media formats is created by enumerating the external codecs installed on the client 305. For example, the codecs installed in the operating system of the client 305 are enumerated by the client agent 375, over the network 315, to the server agent 350. In another embodiment, the list of supported media formats is created by enumerating the codecs embedded in the client agent 375. For example, where the client agent 375 is implemented as a software program, the codecs embedded in the software program are enumerated by the client agent 375, over the network 315, to the server agent 350. Alternatively, the client agent 375 creates the list of supported media formats, and informs the server agent 350 of those supported media formats, by enumerating both the external codecs installed on the client 305 and the codecs embedded in the client agent 375.

[0084] In one embodiment, the client agent 375 generates globally unique identifiers ("GUIDs") and associates each GUID with a particular codec. The client agent 375 then transmits the list of generated GUIDs to the server agent 350 to inform the server agent 350 of the media formats supported by the client agent 375. In another embodiment, the client agent 375 transmits a list of four character codes, each four character code being associated with a particular codec, to the server agent 350 to inform the server agent 350 of the media formats supported by the client agent 375.

[0085] Upon receiving the list of supported media formats from the client agent 375, the server agent 350 registers, at step 408, the first output filter module 355A and/or the second output filter module 355B on the server 310, as associated with the supported media formats. The server agent 350, at step 412, then reports back to the client agent 375 that the server 310 can handle all of the enumerated media formats.

- 21 -

[0086] At step 416, an application program 340 starts executing on the server 310. When the application program 340 identifies within its output, at step 420, the presence of media content, such as, for example, a media stream, the first output filter module 355A, the second output filter module 355B, or both are invoked. If, at step 424, the application program 340 uses external codecs, both the first output filter module 355A and the second output filter module 355B are invoked at step 428 as the application program 340 attempts to invoke an external codec. The first output filter module 355A then intercepts, at step 432, an original compressed data set representing at least a portion of the media stream and transmits, at step 436, the original compressed data set to the client agent 375, without decompressing the data set. The client agent 375, at step 440, receives the original compressed data set and decompresses, at step 444, the original compressed data set to generate a decompressed data set. The client agent 375 uses either external codecs installed on the client 305 or codecs embedded in the client agent 375 itself to decompress the original compressed data set.

[0087] If, instead, at step 424, the application program 340 uses codecs embedded in the application program 340 itself, the second output filter module 355B is, at step 448, invoked as the application program 340 attempts to invoke an OS-level renderer to display the decompressed data set. The second output filter module 355B then intercepts, at step 452, a first decompressed data set, representing at least a portion of the media stream, from the output of the application program 340 and compresses, at step 456, the intercepted first decompressed data set. A variety of compression techniques, including both lossy compression techniques and lossless compression techniques, may be used by the second output filter module 355B, at step 456, to compress the media stream.

[0088] Where the media stream is a video stream, the intercepted first decompressed data set may be compressed, at step 456, by the second output filter module 355B using, for example, a lightweight lossy video encoding algorithm, such as, for example, MJPEG compression. In using the lightweight lossy video encoding algorithm, the second output filter module 355B may choose the desired compression ratio or it may use a predetermined compression ratio. The degree of quality loss chosen by the second output filter module 355B will, typically, depend on the available bandwidth of the network connection. For example, where a user of the client 305 uses a slow modem to connect to the network 315, the second output filter module 355B may choose to use low quality video. Where, on the other hand, a user of the client 305 uses a LAN link or a broadband connection to connect to the network 315, the second output filter module 355B may choose to use a higher quality video.

- 22 -

[0089] Following compression of the intercepted first decompressed data set at step 456, the second output filter module 355B transmits, at step 460, the compressed data set to the client agent 375 in place of the first decompressed data set. The client agent 375, at step 464, receives the compressed data set and decompresses, at step 468, the compressed data set to generate a second decompressed data set. Again, the client agent 375 uses either external codecs installed on the client 305 or codecs embedded in the client agent 375 itself to decompress the compressed data set.

[0090] Regardless of whether the application program 340 uses external or embedded codecs, where the media stream, at step 470, includes a video stream, the second output filter module 355B, at step 472, captures information for locating images of the video stream on a display screen and transmits the captured information over the network 315 to the client agent 375. The client agent 375, at step 474, receives the information for locating the images of the video stream on the display screen.

[0091] Regardless, again, of whether the application program 340 uses external or embedded codecs and regardless, moreover, of whether the media stream, at step 470, includes a video stream, the server agent 350, at step 476, intercepts and transmits, over the network 315, graphical display commands, associated with the non-media graphical information outputted by the application program 340, to the client agent 375. The client agent 375, at step 480, receives the graphical display commands associated with the non-media graphical information.

[0092] Where, at step 484, the output filter module 355A, 355B captures timing information associated with the media stream, the output filter module 355A, 355B transmits, at step 488, the timing information to the client agent 375. The client agent 375 receives, at step 492, the timing information and generates, at step 496, the media presentation at the presentation interface 345. To generate the media presentation at the presentation interface 345, the client agent 375 uses the timing information, the graphical display commands associated with the non-media graphical information, and, where the media stream includes a video stream, the information for locating the images of the video stream on a display screen to seamlessly combine the decompressed data set (or, more specifically, where the application program 340 uses embedded codecs, the second decompressed data set) with the non-media graphical information.

[0093] Where the output filter module 355A, 355B does not capture timing information associated with the media stream, the client agent 375 generates, at step 496, the media presentation at the presentation interface 345 using only the decompressed data set (or, more specifically, where the application program 340 uses embedded codecs, the second

- 23 -

decompressed data set), the graphical display commands associated with the non-media graphical information, and, where the media stream includes a video stream, the information for locating the images of the video stream on a display screen .

[0094] The present invention may be provided as one or more computer-readable programs embodied on or in one or more articles of manufacture. The article of manufacture may be a floppy disk, a hard disk, a CD ROM, a flash memory card, a PROM, a RAM, a ROM, or a magnetic tape. In general, the computer-readable programs may be implemented in any programming language. Some examples of languages that can be used include C, C++, or JAVA. The software programs may be stored on or in one or more articles of manufacture as object code.

Equivalents

[0095] The invention can be embodied in other specific forms without departing from the spirit or essential characteristics thereof. In particular, it is understood that a single server may implement both the invention of FIGS. 1 & 2 and the invention of FIGS. 3, 4A, 4B, & 4C. The foregoing embodiments are therefore to be considered in all respects illustrative rather than limiting on the invention described herein. Scope of the invention is thus indicated by the appended claims rather than by the foregoing description, and all changes which come within the meaning and range of equivalency of the claims are therefore intended to be embraced therein.